Secure Inter-Domain Routing
Internet-Draft
Intended status: Informational
Expires: January 5, 2015

K. Sriram
D. Montgomery
US NIST

July 4, 2014

Enhancement to BGPSEC for Protection against Route Leaks
draft-sriram-route-leak-protection-00

Abstract

This document enumerates different types of route leaks based on observed events on the Internet. It illustrates how BGPSEC in its current form (as described in draft-ietf-sidr-bgpsec-protocol-09) already provides protection against all but one of these route-leaks scenarios. The document further discusses a design enhancement to the BGPSEC protocol that will extend protection against this one remaining type of route-leak attack as well. With the inclusion of this enhancement, BGPSEC is expected to provide protection against all types of route-leaks. The document also includes a stopgap method for detection and mitigation of route leaks for the phase when BGPSEC (path validation) is not yet deployed but only origin validation is deployed.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The BGPSEC protocol [I-D.ietf-sidr-bgpsec-protocol] provides cryptographic protection for some aspects of BGP update messages. It offers mechanisms against mis-originations and hijacks of IP prefixes as well as MIMT (man-in-the-middle) AS path modifications. Route leaks [Cowie2013][Cowie2010][Paseka][LRL][Khare] are an additional type of vulnerability in the global BGP routing system against which BGPSEC so far offers only partial protection. In Section 2, different types of vulnerabilities are enumerated based on observed events on the Internet that have been widely regarded as route leaks. This document illustrates how BGPSEC in its current form (as described in [I-D.ietf-sidr-bgpsec-protocol]) already provides protection against all but one of these route-leaks scenarios. The document further discusses a design enhancement to the BGPSEC protocol that will extend protection against this one remaining type of route-leak attack as well. With the inclusion of this enhancement, BGPSEC is expected to provide protection against all types of route-leaks. The document also presents a stopgap method for detection and mitigation of route leaks for the phase when BGPSEC (path validation) is not yet deployed but only origin validation is deployed.

2. Classification of Route Leaks Based on Documented Events

We use the same basic definition of route leaks here as in [I-D.ietf-grow-simple-leak-attack-bgpsec-no-help]. As illustrated in Figure 1, a route leak occurs when a multi-homed customer AS (such as AS1 in Figure 1) learns a prefix update from one provider (ISP1) and leaks the update to another provider (ISP2) in violation of expected routing policies, and further the second provider does not detect the leak and propagates the leaked update to its customers, peers, and transit ISPs.

```
                                      /\              /\
                                       \ route-leak(P)/
                                        \ propagated /
                                         \          /
              +------------+    peer    +------------+
        ______| ISP1 (AS2) |----------->|  ISP2 (AS3)|---------->
       /       ------------+  prefix(P) +------------+ route-leak(P)
      | prefix |          \   update      /\        \  propagated
       \  (P)  /           \              /          \
        -------   prefix(P) \            /            \
                     update  \          /              \
                              \        /route-leak(P)  \/
                              \/      /
                           +---------------+
                           | customer(AS1) |
                           +---------------+
```

Figure 1: Illustration of the basic notion of a route leak.

Different types of route leaks can be enumerated as follows based on observed attacks on the Internet that have been widely regarded as route leaks:

- o Type 1 (Prefix Hijack with Data Path to Legitimate Origin): A multi-homed AS learns a prefix route from one upstream ISP and announces the prefix to another upstream ISP as if it is being originated by it (i.e. strips the received AS path, and re-originates the prefix). This amounts to straightforward hijacking. Somehow (not attributable to the use of path poisoning trick by the attacker) a reverse path is present, and data packets reach the legitimate destination albeit via the offending AS. But sometimes the reverse path may not be there, and data packets get dropped when received by the offending AS.

  * Examples of this type of route leak include the Iceland and the Belarus incidents in 2013 [Cowie2013], and the China Telecom incident in April 2010 [Cowie2010] [Labovitz].

- Type 2 (U-Turn with More Specific Prefix): A multi-homed AS learns a prefix route from one upstream ISP and announces a sub-prefix (subsumed in the prefix) to another upstream ISP. The AS path in the update is not altered. Update is crafted by the attacker to have a subprefix to maximize the success of the attack while reverse path is kept open by the path poisoning techniques as in [Kapela-Pilosov]. Data packets reach the legitimate destination albeit via the offending AS.

- Type 3 (U-Turn with Full Prefix): A multi-homed AS learns a prefix route from one upstream ISP and simply propagates the prefix to another upstream ISP. Neither the prefix nor the AS path in the update is altered. This is similar to a straight forward path-poisoning attack [Kapela-Pilosov], but with full prefix. The update basically makes a U-turn at the attacker’s multi-homed AS. The attacker often succeeds because the second ISP prefers customer announcement over peer announcement of the same prefix.

  * Examples of Type 3 route leak are the Moratel announcement of Google prefixes in 2012 [Paseka] and the Dodo-Telstra incident in 2012 [Huston].

- Type 4 (Leak of Internal Prefixes): An offending AS simply leaks its internal prefixes to one or more of its provide ASes. The provider AS fails to filter.

3. Mechanisms for Protection against Route Leaks

It is easy to observe that route leaks of Types 1 and 4 (described in Section 2) can be already detected and mitigated by the RPKI-based origin validation alone. This is because, in the case of Type 1 and Type 4, there would be no ROAs available to validate a re-originated prefix or subprefix, and hence the update will be considered Invalid. Assuming that BGPSEC is in use, in the case of a Type 1 route leak, the update will be Invalid due to Invalid path signatures as well as Invalid origin AS. Now turning our attention to route leaks of Type 2, they can be detected and mitigated already by BGPSEC. This is because, in the case of Type 2, changing a prefix to a subprefix (i.e. more specific) in BGPSEC will amount to update modification by an MITM (even though AS path is not modified), and hence the path signatures in the update would no longer be Valid. So the only remaining type of route leaks that needs to be addressed is Type 3.

In Section 3.1 and Section 3.2, we describe a simple addition to BGPSEC that facilitates cryptographically-enabled detection of the Type 3 route leaks as well. Thus, with this enhancement, BGPSEC will be capable of providing detection and mitigation capability against all the different types of route leaks discussed in Section 2.

3.1. Route Leak Protection (RLP) Field Encoding by Sending Router

The key principle is that, in the event of a route leak, a receiving router in a provider AS (e.g. referring to Figure 1, ISP2 (AS3) router) should be able to detect from the prefix-update that its customer AS (e.g. AS1 in Figure 1) SHOULD NOT have forwarded the update (towards the provider AS). This means that at least one of the ASes in the AS path of the update has indicated that it sent the update to its customer or peer AS, but forbade any subsequent ’Up’ forwarding (i.e. from a customer AS to its provider AS). For this purpose, a Route Leak Protection (RLP) field to be set by a sending router is proposed to be used for each AS hop.

For the purpose of route leak detection and mitigation proposed in this document, the RLP field value SHOULD be set to one of two values as follows:

o 00: This is the default value (i.e. "nothing specified"),

o 01: This is the ’Do not Propagate Up’ indication; sender indicating that the prefix-update SHOULD NOT be forwarded ’Up’ towards a provider AS.

There are two different scenarios when a sending AS SHOULD set the ’01’ indication in a prefix-update: (1) when sending the prefix-update to a customer AS, and (2) to let a peer AS know not to forward the prefix-update ’Up’ towards a provide AS. In essence, in both scenarios, the intent of ’01’ indication is that any receiving AS along the subsequent AS path SHOULD NOT forward the prefix-update ’Up’ towards its (receiving AS’s) provider AS.

One may argue for an RLP field value (e.g. ’10’) to be used to specify ’Up’ (i.e. towards provider AS) directionality. But in the interest of keeping the methodology simple, the choice of two RLP field values as defined above (00 - default, and 01 - ’Do not Propagate Up’) is all that is needed. This two-state specification in the RLP field can be shown to work for detection and mitigation of route leaks of Type 3, which is the focus here. (Please see Section 5 for further discussion about the downside using ’Up’ indication.)

The RLP field can be incorporated within the Flags field in the Secure_Path Segment in BPGSEC updates [I-D.ietf-sidr-bgpsec-protocol]. The Flags field in BGPSEC in one octet long, and one Flags field is available for each AS hop, and currently only the first bit is used in BGPSEC. So there are 7 bits that are currently unused in the Flags field. Two of these bits can be designated for the RLP field.

The BGPSEC protocol is expected to provide cryptographic protection to the RLP field. Since the BGPSEC protocol specification requires a sending AS to include the Flags field in the data that are signed over, the RLP field for each hop (assuming it would be part of the Flags field) will be protected under the sending AS’s signature.

3.2. Recommended Actions at a Receiving Router

We provide here an example set of receiver actions that work to detect and mitigate route leaks of Type 3 (in particular). This example algorithm serves as a proof of concept. However, other receiver algorithms or procedures can be designed (based on the same sender specification as in Section 3.1) and may perform with greater efficacy, and are by no means excluded.

A recommended receiver algorithm for detecting a route leak is as follows:

A receiving BGPSEC router SHOULD mark an update as a Route-Leak if ALL of the following conditions hold true:

1. The update is received from a customer AS.

2. It is Valid in accordance with the BGPSEC protocol.

3. The update has the RLP field set to ’01’ (i.e. ’Do not Propagate Up’) indication for one or more hops (excluding the most recent) in the AS path.

The reason for stating "excluding the most recent" in the above algorithm is as follows. The provider AS already knows that most recent hop in the update is from its customer AS to itself, and hence it does not need to rely on the RPL field value set by the customer for detection of route leaks. (See further discussion in Section 5.1.)

After applying the above detection algorithm, a receiving router may use any policy-based algorithm of its own choosing to mitigate any detected route leaks. An example receiver algorithm for mitigating a route leak is as follows:

- If an update from a customer AS is marked as a Route-Leak, then the receiving router SHOULD prefer a Valid signed update from a peer or an upstream provider over the customer's update.

The basic principle here is that the presence of '01' value in the RLP field corresponding to one or more AS hops in the AS path of an update coming from a customer AS informs a provider AS that a route leak is likely occurring. The provider AS then overrides the "prefer customer route" policy, and instead prefers a route learned from a peer or another upstream provider over the customer's route.

A receiving router expects the RLP field value for any hop in the AS path to be either 00 or 01. However, if a different value (say, 10 or 11) is found in the RLP field, then an error condition will get flagged, and any further action is TBD.

4. Stopgap Solution when Only Origin Validation is Deployed

During a phase when BGPSEC has not yet been deployed but only origin validation has been deployed, it would be good have a stopgap solution for route leaks. The stopgap solution can be in the form of construction of a prefix filter list from ROAs. A suggested procedure for constructing such a list comprises of the following steps:

- ISP makes a list of all the ASes (Cust_AS_List) that are in its customer cone (ISP's own AS is also included in the list). (Some of the ASes in Cust_AS_List may be multi-homed to another ISP and that is OK.)

- ISP downloads from the RPKI repositories a complete list (Cust_ROA_List) of valid ROAs that contain any of the ASes in Cust_AS_List.

- ISP creates a list of all the prefixes (Cust_Prfx_List) that are contained in any of the ROAs in Cust_ROA_List.

- Cust_Prfx_List is the allowed list of prefixes that are permitted by the ISP's AS, and will be forwarded by the ISP to upstream ISPs, customers, and peers.

- Any prefix not in Cust_Prfx_List but announced by any of the ISP's customers is marked as a potential route leak. Then the ISP's router SHOULD prefer an Valid (i.e. valid according to origin validation) update from a peer or an upstream provider over the customer’s update for that prefix.

Special considerations with regard to the above procedure may be needed for DDoS mitigation service providers. They typically originate or announce a DDoS victim’s prefix to their own ISP on a short notice during a DDoS emergency. Some provisions would need to be made for such cases, and they can be determined with the help of inputs from DDoS mitigation service providers.

5. Design Rationale and Discussion

In this section, we will try to provide design justifications for the methodology specified in Section 3, and also answer some anticipated questions.

5.1. Downside of ’Up (Towards Provider AS)’ Indication in the RLP Field

As we have shown in Section 3, route leak detection and mitigation can be performed without the use of ’Up’ (i.e. from customer AS to provider AS) indication in the RLP field. The detection and mitigation action should primary occur at a provider AS’s router just as soon as a leaked update is received from a customer AS. At that point, a provider AS can be fooled if it merely looks to see if an offending customer AS has set an ’Up’ indication in the RLP field. This is so since a customer AS intent on leaking a route can deliberately set "Not Specified (00)" indication in order to misguide its provider AS. So it seems better that a provider AS figures out that the update is moving in the ’Up’ direction based only on its own (configuration-based) knowledge that the update is coming from one of its customer ASes. An ’Up’ indication (if it were allowed) can be also potentially misused. For example, an AS in the middle can determine that a ’01’ (i.e. ’Do not Propagate Up’) value already exists on one of the preceding AS hops in a received update’s AS path. Then, said AS in the middle can deliberately set its own RLP field to signal ’Up’, in which case the update may be erroneously marked as a route leak by a subsequent AS if it concludes that there was a valley in the AS path of the update. So there appears to be some possibility of misuse of ’Up’ indication, and hence we proposed not including it in the RLP specification in Section 3. However, other proposals, if any, that aim to beneficially use an ’Up’ indication in the RLP field would be worth discussing.

5.2. Any Possibility of Abuse of ’01’ (i.e. ’Do not Propagate Up’) Indication in the RLP Field?

In reality, there appears to be no gain or incentive for an AS to falsely set its own RLP field to ’01’ (i.e. ’Do not Propagate Up’) indication in an update that it originates or forwards. The purpose of a deliberate route leak by an AS is to attract traffic towards itself, but if the AS were to falsely set its own RLP field to ’01’ value, it would be effectively repelling traffic away from itself for the prefix in question (see receiver algorithm in Section 3.2).

5.3. Route Leaks that Have to Do with Three or More Very Large ISP ASNs in a Sequence in the AS Path

In [Mauch-nanog][Mauch], route leaks of a different kind are characterized by finding three or more very large ISP ASes in a sequence in a BGP update’s AS path. Mauch observes that these are anomalies and potentially route leaks because very large ISPs such as ATT, Sprint, Verizon, Globalcrossing, etc. do not in general buy transit services from each other. However, he also notes that there are exceptions when one very large ISP does indeed buy transit from another very large ISP, and he excludes known cases from his detection algorithm. Because of these exceptions, it is not possible to have a formal definition for the type of route leaks that [Mauch] reports. It may also be noted that route leaks of this type do happen very frequently [Mauch]. Even though they do not seem to generate news in the trade press, they do exist and are a cause for concern. We are keen to develop a better understanding of this topic, and explore additional solution mechanisms that could help detect and mitigate this type of route leaks as well.

6. Security Considerations

The proposed Route Leak Protection (RLP) field requires cryptographic protection. Since it is proposed that the RLP field be included in the Flags field in the Secure_Path Segment in BPGSEC updates, the cryptographic security mechanisms in BGPSEC are expected to also apply to the RLP field. The reader is therefore directed to the security considerations provided in [I-D.ietf-sidr-bgpsec-protocol].

7. IANA Considerations

No updates to the registries are suggested by this document.

8. Acknowledgements

Thanks are due to Danny McPherson for email communication relating to the idea of construction of customer prefix filters using RPKI ROA information. The authors are also thankful to Oliver Borchert and Okhee Kim for their comments.

9. References

9.1. Normative References

[I-D.ietf-sidr-bgpsec-protocol]
Lepinski, M., "BGPSEC Protocol Specification", draft-ietf-sidr-bgpsec-protocol-09 (work in progress), July 2014.

9.2. Informative References

[Cowie2010]
Cowie, J., "China’s 18 Minute Mystery", Renesys Blog, November 2010, <http://www.renesys.com/2010/11/chinas-18-minute-mystery/>.

[Cowie2013]
Cowie, J., "The New Threat: Targeted Internet Traffic Misdirection", Renesys Blog, November 2013, <http://www.renesys.com/2013/11/mitm-internet-hijacking/>.

[Huston] Huston, G., "Leaking Routes", March 2012, <http://labs.apnic.net/blabs/?p=139/>.

[I-D.ietf-grow-simple-leak-attack-bgpsec-no-help]
McPherson, D., Amante, S., Osterweil, E., and D. Mitchell, "Route-Leaks & MITM Attacks Against BGPSEC", draft-ietf-grow-simple-leak-attack-bgpsec-no-help-04 (work in progress), April 2014.

[Kapela-Pilosov]
Pilosov, A. and T. Kapela, "Stealing the Internet: An Internet-Scale Man in the Middle Attack", DEFCON-16 Las Vegas, NV, USA, August 2008, <https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf/>.

[Khare] Khare, V., Ju, Q., and B. Zhang, "Concurrent Prefix Hijacks: Occurrence and Impacts", IMC 2012, Boston, MA, November 2012, <http://www.cs.arizona.edu/~bzhang/paper/12-imc-hijack.pdf/>.

[LRL] Khare, V., Ju, Q., and B. Zhang, "Large Route Leaks", Project web page, 2012, <http://nrl.cs.arizona.edu/projects/lsrl-events-from-2003-to-2009/>.

[Labovitz] Labovitz, C., "Additional Discussion of the April China BGP Hijack Inciden", Arbor Networks IT Security Blog, November 2010, <http://www.arbornetworks.com/asert/2010/11/additional-discussion-of-the-april-china-bgp-hijack-incident/>.

[Mauch] Mauch, J., "BGP Routing Leak Detection System", Project web page, 2014, <http://puck.nether.net/bgp/leakinfo.cgi/>.

[Mauch-nanog] Mauch, J., "Detecting Routing Leaks by Counting", NANOG-41 Albuquerque, NM, USA, October 2007, <https://www.nanog.org/meetings/nanog41/presentations/mauch-lightning.pdf/>.

[Paseka] Paseka, T., "Why Google Went Offline Today and a Bit about How the Internet Works", CloudFare Blog, November 2012, <http://blog.cloudflare.com/why-google-went-offline-today-and-a-bit-about/>.

Authors' Addresses

Kotikalapudi Sriram
US NIST

Email: ksriram@nist.gov

Doug Montgomery
US NIST

Email: dougm@nist.gov